

Application for United States Patent

Title: Network Fault Alerting System and Method

Inventors: Mohammad Ahmed (Citizenship: USA)

8701 Bluffstone Cove, Apt. 7103

5 Austin, Texas 78759

Agent: Robert H. Frantz

Registration Number 42,553

P.O. Box 23324

10 Oklahoma City, OK 73123-2324

CROSS-REFERENCE TO RELATED APPLICATIONS

(CLAIMING BENEFIT UNDER 35 U.S.C. 120)

None.

FEDERALLY SPONSORED RESEARCH

5

AND DEVELOPMENT STATEMENT

This invention was not developed in conjunction with any Federally sponsored contract.

MICROFICHE APPENDIX

Not applicable.

10

Network Fault Alerting System and Method

BACKGROUND OF THE INVENTION

Field of the Invention

This invention pertains to the arts of computer network management, and especially to the management of network bandwidth consumed by network management, status, and maintenance messages. More particularly, this invention relates to the arts of intelligent processing and diagnosis of network failures and problems based on fault analysis logic to more accurately detect and isolate computer network problems, to minimize the network bandwidth consumed by maintenance messages, and to effectively notify maintenance personnel of the most likely point of failure.

Description of the Related Art

Computer networks, such as local area networks ("LAN"), wide-area networks ("WAN"), intranets and the Internet typically include substantial maintenance and monitoring capabilities. Modern telephone networks, such as Signalling System 7 ("SS7), Integrated Services Data Network ("ISDN"), and many digital cellular networks including GSM, also include substantial equipment and software which are dedicated to the provisioning, monitoring and maintenance of the network and its equipment. All of the above named networks are packet-based networks, and are well-known within their respective arts.

Key to the architecture and operation of these networks are packet routers, which interconnect multiple physical networks and provide routing and forwarding of packets, or "messages", from one network to another based upon addressing schemes defined by well-known protocols such as the Internet Protocol ("IP") or LAPD for SS7 and ISDN. These addressing schemes can be generalized as schemes which define each data packet or message has having a header, payload, and tail. The destination address, origination address, packet sequence number, and payload size are typically included in the header section of the message. The payload section contains the actual computer data which is being transferred from one computer to another via the computer network, which may represent a portion of a computer file, a formatted message, or a section of digitized signal such as voice, video or other audio. The various message formats are defined by well-known standards promulgated by InterNIC, the International Telecommunications Union, Bellcore, and ANSI.

In order to manage these networks, including monitoring of network operation status, configuring and re-configuring network elements (routers, terminals and switches), and provisioning of new network sections, a number of well-known software and hardware products have been developed and placed on the market. Most of these products integrate specialized software onto network server platforms. The software uses the network connectivity and bandwidth provided by the network server platform to perform maintenance testing, messaging, status checking, and alert messaging. Many times, the actual network being used for "real" traffic, such as computer file transmission or telephone call transmission, is used for the maintenance

communications as well. In this case, the maintenance messages “mix in” with the bandwidth of the “real” traffic. As such, if maintenance messages accumulate to significant bandwidth consumption, network performance may be adversely affected.

In other cases, separate networks dedicated to maintenance may be configured to

- 5 avoid this problem. But, even so, if maintenance messages exceed an expected bandwidth level, the dedicated maintenance network may fail.

When network management software like Netview/6000 or Hewlett-Packard's OpenView and others, detects a network device such as a router has gone off-line, it will send “node down” events or messages for all the workstations connected

- 10 downstream from off-line router to network problem management server. The network problem management server provides correlation and processing for opening trouble tickets, and eventually, it send alerts to appropriate maintenance personnel thru pagers, e-mail, and/or telephone calls.

- FIGURE 1 shows the topology of prior art maintenance systems. A router (1) may have multiple ports to multiple networks. Each port is serviced by a network interface card (“NIC”), such as an Ethernet LAN interface card. FIGURE 1 shows an example of a router serving three networks, A, B, and C, each of which is a group of networked computer workstations or personal computers. For example, network A (5) has several “drops” to computers, and one drop or connection (6) to the router. Likewise, network B (4) is connected (3) to the router, and network C (2) is connected (7) to the router. Packets or messages received by the router are forwarded to other

networks based on the addressing scheme of the network, such as IP in the case of many computer networks.

Also shown in FIGURE 1 is a connection (8) to a maintenance server (9) such as a NetView 6000 server. In this example, this connection (8) connects to the router (1) using the router's NIC for network D. The maintenance server (9) typically contains a connectivity database which contains all of the network addresses of all the elements on the other networks connected to the router, such as all the computers connected to networks A, B, and C. Using this database, the maintenance server (8) periodically sends status query messages, or "pings", to each of the computers. If each computer is on-line, the router is functioning properly, and the network physical media (cable, RF links, etc.) is in tact, a reply will be received from each computer nearly immediately in response to the "ping". If a reply or response is not received within a certain time from transmitting of the "ping", the maintenance server (9) may assume a problem with the computer, router, or network(s) exists.

For example, if all computers and the router are functioning correctly except for one computer, then only one response will not be received, and all other responses will be received. However, if the router fails, no responses will be received from any of the computers. In the most basic of maintenance system configurations such as the basic NetView 6000 product, this scenario can result in a storm of events being sent to the problem management server which correlates events and opens trouble tickets, leading to many useless and/or redundant e-mails and pagers.

FIGURE 2 illustrates this scenario. A normal "ping" (20) is forwarded from the NetView 6000 to the router, which forwards (21) it to the appropriate PC. The PC, if functioning properly, replies (22) via the router to the NetView 6000 (23) within a predetermined time limit t_1 . If the router has failed, the "ping" (24) will not be replied to by any of the computers within time t_1 , which will result in the NetView 6000 sending multiple "computer down" messages (25) to the problem management server. The problem management server is configured to wait a period of time t_3 before escalating the event to notification of the maintenance personnel, in order to reduce the number of alerts made for temporary problems such as power glitches, computer reboots, etc. But, if no "computer up" messages are received within time limit t_3 , the problem management server will send multiple pager messages and telephone calls, and may open multiple trouble tickets (26), as many as one per computer on the network. This results in the alerting of the maintenance personnel, but is confusing to the personnel as to which element is actually failed. Additionally, the network link between the NetView 6000 server and the problem management server has suffered unnecessary bandwidth consumption by all of the "computer down" messages.

In an enhancement of the prior art network management technology, a product called Tivoli for Network Connectivity module (TFNC) by International Business Machines ("IBM") employs similar concept, but it adds some intelligent processing to the maintenance server. With TFNC, all of the original "computer down" messages will be sent to the problem management server, but, as shown in FIGURE 3, the

Tivoli processing (30) will examine the network topology and determine that all of these failures are likely due to a single point failure, namely a router failure. So, within the escalation time period t_3 , TFNC will send multiple "computer up" messages (31) to the problem management server, which results in a net status of only

5 the "router down" message being escalated by the problem management server.

While this enhancement to the network maintenance technology produces a desirable reduction in the number of alerts (pager messages, trouble tickets, etc.) (32) issued to maintenance personnel, it does not reduce the bandwidth consumed by the messages on the network between the maintenance server (TFNC and NetView 6000). Rather,

10 it nearly doubles the bandwidth consumption.

Therefore, there is a need in the art for a system and method which intelligently processing the "ping" response pattern in a timely manner, and which issues a minimal number of "network element down" messages which precisely isolate the most likely point of failure in order to minimize network bandwidth

15 consumption, and to minimize redundant and incorrect maintenance alerts.

BRIEF DESCRIPTION OF THE DRAWINGS

The following detailed description when taken in conjunction with the figures presented herein present a complete description of the present invention.

5 FIGURE 1 shows the prior art topology for network management servers, software, and connectivity.

FIGURE 2 discloses the message sequence used in prior art network management technology.

10 FIGURE 3 discloses the enhanced prior art network management technology message sequence.

FIGURE 4 illustrates the functional flow of the inventive method which filters and diagnosis the most likely point of failure in the network.

FIGURE 5 shows the modified network topology to include a system which implements the inventive method.

15 FIGURE 6 shows the message sequence achieved by use of the inventive method, with substantially reduced network bandwidth requirements and increased accuracy of the alerts.

SUMMARY OF THE INVENTION

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of a preferred embodiment of the invention, as illustrated in the accompanying drawings wherein like reference
5 numbers represent like parts of the invention.

The inventive method is preferably implemented as a software application which will integrate with the existing network management software packages and servers, such as Netview/6000, Hewlett-Packard OpenView, and others. The new software application implements the following general method or logic:

10 (a) When a router or a network device goes off-line, then it will send only one "network element or router down" event to the problem management server which does the correlation and issues the trouble tickets for alerting. Thus when the router down (network device) event is sent via a pager or email, the network operations personnel will know the router is down, and it is obvious that the devices connected
15 downstream will be offline from the entire network;

(b) When a router NIC, port or interface goes off-line, the same logic should result in only one router down message being sent to the problem management server;
and

(c) When a networked element other than a router or NIC, such as a computer,
20 goes off-line, it will send only one "computer down" event to the problem management server.

DETAILED DESCRIPTION OF THE INVENTION

The inventive method is preferably realized as a software application, called "Valerie", which integrates with existing network management software packages and servers, such as Netview/6000, Hewlett-Packard OpenView, on common network server computer hardware platforms such as an IBM RS/6000.

By the logic of the method, it is assumed that it is not likely that multiple failures occur on the same network simultaneously. An even if multiple failures are detected or indicated, certain patterns to the indications allow for diagnosis of a more likely single point of failure. For example, if all but one of the computers on network A in FIGURE 1 are responding to "pings", it is more likely that the non-responsive computer is the failure point as the network wiring, router NIC, and router are still functioning for the other computers on network A. In fact, if even one computer on the network responds, it can be assumed that the network wiring, NIC and router are functioning correctly. However, if the pattern of non-responses includes all of the computers on a network, then the NIC and the router are suspect.

So, in the second step of the logic, if any computers on any other network connected to the router are responding, but all of the computers on just one network are not responding, it can be assumed to be a network wiring or NIC problem with the non-responding network. But, if no computers on any networks are responding, then the router can be assumed to be the single-point of failure.

In order to process the non-responses and the responses in this logical fashion, the Valerie application must have access to the connectivity database which describes

the topology of the networks and computers interconnected by the router, and contains the addresses of the computers and other network elements. This database is already available from or through the network maintenance server, typically through a application program interface ("API"). In the preferred embodiment, Valerie is a

5 software application written in "C", and compiled and targeted for an RS/6000 computer platform running under the AIX operating system concurrently with NetView/6000. However, other languages, such as Java or C++, platforms, such as a Sun Server or IBM-compatible personal computer, and operating systems, such as Solaris or Microsoft Windows NT, may be used as the target system. In any case, the

10 Valerie application program can access the connectivity database via an API through the NetView or OpenView application. Valerie can also send and receive messages using the platform's communication protocol stack, such as IP, and network interface cards, such as Ethernet interfaces, as well as monitor for messages on the network. The integration of Valerie into the overall network management technology is

15 completed by disabling the "element down/element up" message output capability of the NetView or OpenView software, and by enabling the output of the Valerie logic results. Valerie's logic can be triggered by the results of the monitoring activity, or more actively by "trapping" the output event from the NetView or OpenView software.

20 FIGURE 4 summarizes the logic of Valerie in a functional flow depiction. When Valerie is started (41), it reads the connectivity database and develops rules based on the network connectivity related to the router. Then, it periodically sends

“ping” messages (43) to each element connected to the router. Alternatively, it may simply monitor the network for “pings” from the NetView application to each network element. These “pings” can be sent at any interval rate, but are sent at approximately 5 minute intervals in the preferred embodiment. Until a response is not received within a determined time limit, such as 5 minutes, the period “pings” continue. But when one or more responses are not received within the time limit, the logic processing begins. First, a recent history log is examined (44) to determine if any other computers on the same network or router NIC have been received. If so, then a single “element down” message for the non-responding element or computer is sent (45) to the problem management server.

If no other responses have been received recently from other elements or computers on the same network, then the history log is examined (46) to see if any other computers or elements on any other networks connected to the router have been received. If so, then the router NIC and/or network cabling for the the non-responsive network is assumed to be the point of failure, and a single “NIC or network down” message is sent (47) to the problem management server.

However, if no other elements or computers on any of the networks connected to the router have responded recently, then a single “router down” message is sent (48) to the problem management server.

In this embodiment of the invention, the history log can be built and updated by Valerie actively transmitting “pings” to network elements and registering the received responses. Or, it can be built passively by Valerie monitoring (or

"snooping") the network for "pings" and responses between network elements and the network management software application (NetView/6000 or OpenView).

In an alternate embodiment of the invention, the history log is updated by quickly issuing new "pings" to all other network elements when a single response is not received within the time limit. This allows the fault deduction logic to operate on more recent data, giving a more accurate result.

Finally, turning to FIGURE 6, the reduced message bandwidth realized by the invention is noticable. Following the Valeria processing (62), a single "element down" message is sent to the problem management server by the enhanced maintenance server, shown here as NetView/6000 with Valerie.

It will be understood from the foregoing description that various modifications and changes may be made in the preferred embodiment of the present invention without departing from its true spirit, such as the use of alternate programming methodologies or languages, alternate server platforms, various networking protocols, operating systems and development tool sets. It is intended that this description is for purposes of illustration only and should not be construed in a limiting sense. The scope of this invention should be limited only by the language of the following claims.